Object Depth Profile and Reflectivity Restoration from Sparse Single-Photon Data Acquired in Underwater Environments

Abderrahim Halimi, Aurora Maccarone, Aongus McCarthy, Steve McLaughlin and Gerald S. Buller

arXiv:1608.06143v1 [stat.CO] 22 Aug 2016

Abstract

This paper presents two new algorithms for the joint restoration of depth and reflectivity (DR) images constructed from time-correlated single-photon counting (TCSPC) measurements. Two extreme cases are considered: (i) a reduced acquisition time that leads to very low photon counts and (ii) a highly attenuating environment (such as a turbid medium) which makes the reflectivity estimation more difficult at increasing range. Adopting a Bayesian approach, the Poisson distributed observations are combined with prior distributions about the parameters of interest, to build the joint posterior distribution. More precisely, two Markov random field (MRF) priors enforcing spatial correlations are assigned to the DR images. Under some justified assumptions, the restoration problem (regularized likelihood) reduces to a convex formulation with respect to each of the parameters of interest. This problem is first solved using an adaptive Markov chain Monte Carlo (MCMC) algorithm that approximates the minimum mean square parameter estimators. This algorithm is fully automatic since it adjusts the parameters of the MRFs by maximum marginal likelihood estimation. However, the MCMC-based algorithm exhibits a relatively long computational time. The second algorithm deals with this issue and is based on a coordinate descent algorithm. Results on single-photon depth data from laboratory based underwater measurements demonstrate the benefit of the proposed strategy that improves the quality of the estimated DR images.

Index Terms

Lidar waveform, underwater Lidar, Bayesian estimation, Poisson statistics, image restoration, ADMM, MCMC.

The authors are with the School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh U.K.

This work was supported by the EPSRC Grants EP/J015180/1, EP/N003446/1, EP/M01326X/1, EP/K015338/1 and the DSTL National PhD Scheme.

I. Introduction

Reconstruction of 3-dimensional scenes is a challenging problem encountered in many applications. For a given pixel, the time-of-flight light detection and ranging (Lidar) system achieves this goal by emitting laser pulses and recording the round-trip return time and intensity of the reflected signal [1]. Single-photon Lidar typically uses a high repetition rate pulsed laser source in conjunction with a single-photon detector. The advantages of the single-photon approach are its shot-noise limited sensitivity, and its picosecond temporal response which can achieve millimeter-scale surface-to-surface resolution [2]. In single-photon Lidar, the recorded photon event is stored in a timing histogram which is formed by detecting photons from many laser pulses. The time delay and the amplitude of the histogram are related to the distance and reflectivity of the observed object, respectively, which allows the construction of the 3D scene.

In this paper, we consider a scanning system whose acquisition time is defined by the user and is the same for each pixel, which leads to a deterministic and user-defined overall acquisition duration. Consequently, the number of detected photons can be larger than one for some pixels, whereas other pixels may be empty (i.e. no detected photons). We also assume solid target surfaces fabricated from opaque materials, so that only one reflection is observed in an individual pixel [3]. The study focuses on the following two extreme cases: (i) a reduced data acquisition time and (ii) the use of an extremely attenuating medium [4]. Both cases lead to a reduction in the number of detected photons per pixel, which affects the estimation of depth and target reflectivity. Indeed, taking underwater measurements leads to a severe attenuation of the intensity with respect to (w.r.t.) the target range, which makes the reflectivity estimation difficult. With such challenging scenarios, the measurement can be improved by, for example, increasing the laser power or the data acquisition time [5], [6], however this is not always practicable in a field situation. To use the available sparse photon data most efficiently, the alternative approach is to improve the processing of the acquired signals using signal processing techniques [3], [7]–[9]. The latter approach will be considered here to improve the estimated depth and reflectivity (DR) images for sparse single-photon data.

The first contribution of this paper is the use of a hierarchical Bayesian model associated with the DR images. Using the Poisson distribution of the observed photon counts, and introducing some approximations, lead to a log-concave likelihood distribution w.r.t. each of

the parameters of interest. The resulting likelihood distribution is interesting for two reasons: it allows the use of convex programming algorithms for parameter estimation and it is expressed w.r.t. preliminary estimates of the DR images which avoids the use of cumbersome photon count histograms during the refinement process. Using Markov random fields (MRF), the parameters of interest are assigned prior distributions enforcing a spatial correlation between the pixels. More precisely, the depth image is assigned an MRF distribution equivalent to a total variation (TV) prior [10], [11], while the reflectivity image is assigned a gamma-MRF prior [12]. The likelihood and the prior distribution are then used to build the joint posterior distribution that is used for the parameter estimation.

The second contribution of this paper is the derivation of two estimation algorithms associated with the proposed hierarchical Bayesian model. The first algorithm generates samples distributed according to the posterior using Markov chain Monte Carlo (MCMC) methods (such as the Gibbs sampler, and the Metropolis-Hastings algorithm) [13]. These samples are then used to evaluate the minimum-mean-square-error (MMSE) estimator of the DR images. This approach also allows the estimation of the regularization parameters, (the hyperparameters), associated with the MRF prior using the maximum marginal likelihood approach proposed in [14]. Therefore, the MCMC method is fully automatic in the sense that it does not require the user to tune the model hyperparameters. However, the resulting MCMC-based algorithm has a high computational complexity which can be a significant limitation for real time applications. The second algorithm deals with this limitation and approximates the maximum a posteriori (MAP) estimator by using a coordinate descent algorithm [15], [16]. The latter is used to sequentially update the different parameters to minimize the negative log-posterior, which is convex w.r.t. each parameter. In contrast to the reflectivity image that is updated analytically, the depth image is updated using the alternating direction method of multipliers (ADMM). This algorithm has shown good performance in different fields, both for the estimation quality and the reduced computational cost [9], [17], [18]. The proposed algorithms are complementary and represent useful tools to deal with different user requirements such as a reduced computational cost or an automatic hyperparameter estimation. Results on single-photon depth data acquired from laboratory experiments show the benefit of the proposed strategies that improve the quality of the estimated DR images.

The paper is organized as follows. Section II introduces the observation model associated with the underwater photon counts. The proposed hierarchical Bayesian algorithm for DR restoration is presented in Section III. Section IV introduces the two proposed estimation

algorithms based on stochastic simulation and optimization. Simulation results on synthetic data are reported in Section V. Section VI presents and analyzes results conducted using data acquired by an actual time-of-flight scanning sensor based on TCSPC. Finally, conclusions and future work are reported in Section VII.

II. Observation model

The Lidar observation $\boldsymbol{y}_{i,j,t}$, where $(i,j) \in \{1,\cdots,N_r\} \times \{1,\cdots,N_c\}$, represents the number of photon counts within the tth bin of the pixel (i,j). According to [3], [19], each photon count $\boldsymbol{y}_{i,j,t}$ is assumed to be drawn from the Poisson distribution $\mathcal{P}(.)$ as follows

$$y_{i,j,t} \sim \mathcal{P}\left(s_{i,j,t}\right) \tag{1}$$

where $s_{i,j,t}$ is the average photon counts given by [4]

$$s_{i,j,t} = r_{i,j} e^{-\alpha t_{i,j}} g_0\left(t - t_{i,j}\right) + b_{i,j} \tag{2}$$

and $t_{i,j} \geq 0$ is the position of an object surface at a given range from the sensor (related to the depth), $r_{i,j} \geq 0$ is the reflectivity of the target, $b_{i,j} \geq 0$ is a constant denoting the background and dark photon level, α represents the attenuation factor related to the transmission environment and g_0 denotes the system impulse response assumed to be known from the calibration step. In air, the attenuation factor is $\alpha = 0$ and the model (2) reduces to that studied in [3], [9]. This paper considers the case of transmission under a highly attenuating environment in which $\alpha \geq 0$. In this case, the measured reflected intensity of the objects decreases as a function of their distance to the sensor which is valid for different scenarios such as highly scattering underwater measurements. Indeed, the single-photon depth images can be used underwater to localize objects such as boat wreckage, pipelines, etc. The first objective of this paper is to estimate the target depth and reflectivity images of a target underwater or in any other extremely attenuating environment. The paper second objective deals with the extreme case of a very low photon counts per pixels. Under this scenario, it is possible to have missing pixels which have no received photons, i.e., $\sum_{t=1}^{T} y_{i,j,t} = 0$. These missing pixels bring no information regarding the depth $t_{i,j}$ and reflectivity $r_{i,j}$ and should be considered separately from informative observed pixels as in [20].

III. Hierarchical Bayesian Model

This section introduces a hierarchical Bayesian model for estimating the target distance and reflectivity images of underwater measurements. The Bayesian approach accounts for both

the statistical model associated with the observed data (likelihood) and the prior knowledge about the parameters of interest (prior distributions). This approach is interesting to alleviate the indeterminacy resulting from ill-posed problems and has been successfully applied to Lidar measurements in [3]. More precisely, if $f(\boldsymbol{\Theta})$ denotes the prior distribution assigned to the parameter $\boldsymbol{\Theta}$, the Bayesian approach computes the posterior distribution of $\boldsymbol{\Theta}$ using the Bayes rule

$$f(\boldsymbol{\Theta}|\boldsymbol{Y}) \propto f(\boldsymbol{Y}|\boldsymbol{\Theta}) f(\boldsymbol{\Theta}) \tag{3}$$

where $\propto$ means "proportional to" and $f(\boldsymbol{Y}|\boldsymbol{\Theta})$ is the likelihood of the observation matrix $\boldsymbol{Y}$ gathering all the observed pixels $y_{i,j,t}, \forall i, j, t$. The MMSE and MAP estimators of $\boldsymbol{\Theta}$ can be evaluated by the mean vector and maximum of this posterior. At this point, it is interesting to highlight the link between the Bayesian and optimization perspectives. Indeed, the MAP estimator can also be evaluated by minimizing the cost function obtained as the negative log-posterior function. From an optimization perspective, this cost function is considered as a regularized problem where the data fidelity term (likelihood) is constrained using some regularization terms (prior distributions). The following sections introduce the likelihood and the prior distributions (regularization terms) considered in this paper.

A. Likelihood

Assuming independence between the observed pixels $y_{i,j,t}$ and considering the Poisson statistics leads to the following joint likelihood

$$P(\boldsymbol{Y}|\boldsymbol{t}, \boldsymbol{r}, \boldsymbol{b}) = \prod_{(i,j)\in\Omega} \prod_{t=1}^{T} \frac{s_{i,j,t}^{y_{i,j,t}}}{y_{i,j,t}!} \exp^{-s_{i,j,t}} \tag{4}$$

where $\boldsymbol{t}, \boldsymbol{r}, \boldsymbol{b}$ are $N \times 1$ vectors gathering the elements $t_{i,j}$, $r_{i,j}$, $b_{i,j}, \forall i, \forall j$ (in lexicographic order), with $N = N_r N_c$, T is the total number of bins, Ω gathers the indices of non-empty pixels and $s_{i,j,t}(\boldsymbol{t}, \boldsymbol{r}, \boldsymbol{b})$ has been denoted by $s_{i,j,t}$ for brevity. In a similar fashion to the classical estimation approach (see [3], [9] for more details), this paper assumes the absence of the background level, i.e., $b_{i,j} = 0$. Indeed, the underwater measurements are most often obtained in dark conditions (in the laboratory in our case) which justifies this assumption. Note, however, that the assumption is violated in presence of multiple scatterers, thus, its effect is studied when considering synthetic data. In addition to this simplification, we further assume a Gaussian approximation for the instrument impulse response[1] $g_0(t - t_{i,j}) =$

[1]The parameters c_1 and σ^2 can be estimated by fitting the actual impulse response with a Gaussian using a least squares algorithm.

$c_1 \exp^{-\frac{\left(t-t_{i,j}\right)^2}{2\sigma^2}}$ as in [9], [21], and that the temporal sum of the shifted impulse response $c_2 = \sum_{t=1}^{T} g_0\left(t - t_{i,j}\right)$ is a constant for all realistic target distances $t_{i,j}$ (which is justified when assuming that the observation time window is larger than the depth of the observed object). Under these assumptions, the likelihood reduces to $\mathcal{L} = \prod_{(i,j)\in\Omega} \mathcal{L}_{i,j}$ with (after removing unnecessary constants)

$$\mathcal{L}_{i,j} = r_{i,j}^{c_2 r_{i,j}^{\mathrm{ML0}}} \exp\left[-\alpha c_2 r_{i,j}^{\mathrm{ML0}} t_{i,j} - \frac{\left(t_{i,j} - t_{i,j}^{\mathrm{ML0}}\right)^2}{\frac{2\sigma^2}{c_2 r_{i,j}^{\mathrm{ML0}}}} - c_2 r_{i,j} \exp^{\left(-\alpha t_{i,j}\right)}\right] \tag{5}$$

where $t_{i,j}^{\mathrm{ML0}} = \frac{\left(\sum_{t=1}^{T} t y_{i,j,t}\right)}{\left(\sum_{t=1}^{T} y_{i,j,t}\right)}$ and $r_{i,j}^{\mathrm{ML0}} = \frac{1}{c_2}\left(\sum_{t=1}^{T} y_{i,j,t}\right)$ are the maximum of this simplified likelihood w.r.t. $t_{i,j}$ and $r_{i,j}$ obtained in the air (with $\alpha = 0$). The likelihood (5) obtained is interesting for two reasons. First, it does not include the Lidar observation terms $y_{i,j,t}$ explicitly, which means that our formulation considers only the two observed images $r_{i,j}^{\mathrm{ML0}}$ and $t_{i,j}^{\mathrm{ML0}}$ instead of the $N_r \times N_c \times T$ matrix $y_{i,j,t}$. The computational cost is then drastically reduced when compared to the models studied in [3], [22] which considered the full $N_r \times N_c \times T$ data cube. Second, it is a log-concave distribution w.r.t. each of the parameters $t_{i,j}$ and $r_{i,j}$ separately, that is suitable for the application of convex programming algorithms. Note finally that our approach can be interpreted as a joint depth-reflectivity image restoration problem of the estimates $t_{i,j}^{\mathrm{ML0}}$ and $r_{i,j}^{\mathrm{ML0}}$ that are of poor quality especially in the limit of very low photon counts or when acquiring the data in a significantly attenuating environment. The next section introduces the prior information introduced to improve the estimated images from (5).

B. Priors for the distance image

The target distances exhibit correlation between adjacent pixels. This effect is accounted for by considering the following MRF prior distribution

$$f(\boldsymbol{t}|\eta) = \frac{1}{G(\eta)} \exp^{[-\eta \mathrm{TV}(\boldsymbol{t})]} \tag{6}$$

where $G(\eta)$ is a normalizing constant, η is a coupling parameter that controls the amount of enforced spatial smoothness, $\mathrm{TV}(\boldsymbol{t}) = \sum_{i,j} \sum_{(i',j')\in\upsilon(i,j)} |t_{i,j} - t_{i',j'}|$ denotes the total-variation regularization suitable for edge preservation [10], [11] and $\upsilon(i,j)$ denotes the neighborhood of the pixel (i,j) as shown in Fig. 1.

C. Priors for the reflectivity image

Similarly as for the target distances, we expect the target reflectivity to vary smoothly from one pixel to another. This behavior is obtained by introducing an auxiliary variable $\boldsymbol{w}$ (of size $N_{\mathrm{r}} \times N_{\mathrm{c}}$) and assigning a gamma-MRF prior for $(\boldsymbol{r}, \boldsymbol{w})$ as follows [12], [23], [24]

$$\begin{aligned} f\left(\boldsymbol{w}, \boldsymbol{r}|\zeta\right) = & \quad \tfrac{1}{Z(\zeta)} \prod_{(i,j)\in\nu_{\boldsymbol{w}}} w_{i,j}^{-(4\zeta+1)} \\ \times & \quad \prod_{(i',j')\in\nu_{\boldsymbol{r}}} r_{i',j'}^{(4\zeta-1)} \\ \times & \quad \prod_{((i,j),(i',j'))\in\mathcal{E}} \exp\left(\tfrac{-\zeta r_{i',j'}}{w_{i,j}}\right), \end{aligned} \tag{7}$$

where $Z(\zeta)$ is a normalizing constant, the partition $\nu_{\boldsymbol{w}}$ (resp. $\nu_{\boldsymbol{r}}$) denotes the collection of variables $\boldsymbol{w}$ (resp. $\boldsymbol{r}$), the edge set $\mathcal{E}$ consists of pairs (i,j) representing the connection between the variables and ζ is a coupling parameter that controls the amount of spatial smoothness enforced by the GMRF. This prior ensures that each $r_{i,j}$ is connected to four neighbor elements of $\boldsymbol{w}$ and vice-versa (see Fig. 2). The reflectivity coefficients $r_{i,j}$ are conditionally independent and the 1st order neighbors (i.e., the spatial correlation) is only introduced via the auxiliary variables $\boldsymbol{w}$. An interesting property of this joint prior is that the conditional prior distributions of $\boldsymbol{r}$ and $\boldsymbol{w}$ reduce to conjugate inverse gamma ($\mathcal{IG}$) and gamma ($\mathcal{G}$) distributions as follows

$$\begin{aligned} w_{i,j}|\boldsymbol{r}, \zeta &\sim \mathcal{IG}\left(4\zeta, 4\zeta\rho_{1,i,j}(\boldsymbol{r})\right), \\ r_{i,j}|\boldsymbol{\epsilon}, \zeta &\sim \mathcal{G}\left(4\zeta, 1/(4\zeta\rho_{2,i,j}(\boldsymbol{w}))\right), \end{aligned} \tag{8}$$

where

$$\begin{aligned} \rho_{1,i,j}(\boldsymbol{r}) &= (r_{i,j} + r_{i-1,j} + r_{i,j-1} + r_{i-1,j-1})/4, \\ \rho_{2,i,j}(\boldsymbol{w}) &= (w_{i,j}^{-1} + w_{i+1,j}^{-1} + w_{i,j+1}^{-1} + w_{i+1,j+1}^{-1})/4. \end{aligned} \tag{9}$$

D. Posterior distribution

The proposed Bayesian model is illustrated by the directed acyclic graph (DAG) displayed in Fig. 3, which highlights the relation between the observations $\boldsymbol{Y}$, the parameters $\boldsymbol{t}, \boldsymbol{r}, \boldsymbol{w}$ and the hyperparameters η, ζ. Assuming prior independence between the parameter vector $\boldsymbol{\Theta} = (\boldsymbol{t}, \boldsymbol{r}, \boldsymbol{w})$, the joint posterior distribution associated with the proposed Bayesian model is given by

$$f\left(\boldsymbol{\Theta}|\boldsymbol{Y}, \eta, \zeta\right) \propto f(\boldsymbol{Y}|\boldsymbol{\Theta}) f\left(\boldsymbol{\Theta}|\eta, \zeta\right). \tag{10}$$

This posterior will be used to evaluate the Bayesian estimators of $\boldsymbol{\Theta}$. For this purpose, we propose two algorithms based on an MCMC and an optimization approach. The first approach uses an MCMC approach to evaluate the MMSE estimator of $\boldsymbol{\Theta}$ by generating samples according to the joint posterior distribution. Moreover, it allows the estimation of the hyperparameters η, ζ by using a maximum marginal likelihood estimation during the inference procedure (as detailed in the next section). However, this MCMC algorithm presents a significant computational complexity which can limit the applicability for real time applications. The second optimization algorithm deals with this issue and provides fast MAP estimates for $\boldsymbol{\Theta}$. This is achieved by maximizing the posterior (10) w.r.t. $\boldsymbol{\Theta}$, or equivalently, by minimizing the negative log-posterior given by $\mathcal{F} = -\log[f(\boldsymbol{\Theta}|\boldsymbol{Y}, \eta, \zeta)]$. Note however, that the hyperparameters are fixed under this approach. The two estimation algorithms are described in the next section.

IV. Estimation algorithms

A. MCMC algorithm

The principle of the MCMC approach is to generate samples whose stationary distribution is the desired posterior distribution (10). The distribution (10) being difficult to sample, the Gibbs algorithm can be used to iteratively generate samples according to its conditional distributions [13]. Moreover, when a conditional distribution cannot be sampled directly, sampling techniques such as the Metropolis-Hasting (MH) algorithm can be applied leading to a Metropolis-within-Gibbs sampler. In this paper, we generate samples associated with the parameters $(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3) = (\boldsymbol{t}, \boldsymbol{r}, \boldsymbol{w})$ and use them to approximate the MMSE estimators given by

$$\hat{\boldsymbol{\theta}}_i^{\text{MMSE}} = \mathbb{E}\left[\boldsymbol{\theta}_i|\boldsymbol{Y}, \hat{\eta}, \hat{\zeta}\right], \text{ for } i = 1, 2, 3 \tag{11}$$

where the expectation $\mathbb{E}(.)$ is taken w.r.t. the marginal posterior density $f\left(\boldsymbol{\theta}_i|\boldsymbol{Y}, \hat{\eta}, \hat{\zeta}\right)$ (by marginalizing $\boldsymbol{\theta}_j, j \neq i$, this density takes into account their uncertainty). In addition to these parameters, the hyperparameters η, ζ are also estimated by considering the method proposed in [14], which is based on the maximum marginal likelihood estimator, given by

$$\left(\hat{\eta}, \hat{\zeta}\right) = \underset{\eta \in \mathbb{R}^+, \zeta \in \mathbb{R}^+}{\text{argmax}} \; f(\boldsymbol{Y}|\eta, \zeta). \tag{12}$$

This method provides a point estimate for the hyperparameters that is used to evaluate the parameter MMSE as indicated in (11). These approaches have two main advantages: (i) it

allows for an automatic adjustment of the value of (η, ζ) for each image which leads to an estimation improvement, (ii) it has a reduced computational cost when compared to competing approaches [25]. It should be noted that the resulting algorithm is similar to [3] while the main differences relate to the different estimated parameters, the different distribution expressions (due to the underwater observation model), and to the discrete target positions in [3] which are continuous in this paper. The next subsections provide more details regarding the main steps of the sampling algorithm.

1) Sampling the target positions: The conditional distribution of $\boldsymbol{t}$ is given by

$$f\left(\boldsymbol{t}|\boldsymbol{r},\boldsymbol{Y}\right) \propto \exp^{-\mathcal{C}(\boldsymbol{t})}, \tag{13}$$

with

$$\begin{aligned}\mathcal{C}(\boldsymbol{t}) &= \sum_{(i,j)\in\Omega}\left[\frac{\left(t_{i,j}-t_{i,j}^{\text{ML0}}+\alpha\sigma^2\right)^2}{\frac{2\sigma^2}{c_2 r_{i,j}^{\text{ML0}}}}+c_2 r_{i,j}\exp^{(-\alpha t_{i,j})}\right]\\ &+ \; i_{\mathbb{R}_+}\left(\boldsymbol{t}\right)+\eta\text{TV}\left(\boldsymbol{t}\right)\end{aligned} \tag{14}$$

where the observations $\boldsymbol{Y}$ are introduced via the images $r_{i,j}^{\text{ML0}}$ and $t_{i,j}^{\text{ML0}}$ and $i_{\mathbb{R}_+}\left(\boldsymbol{t}\right)$ is the nonnegative orthant indicator function. Since it is not easy to sample according to (13), we propose to update the target positions using a Metropolis-Hasting (MH) move. More precisely, a new position is proposed following a Gaussian random walk procedure (the variance of the proposal distribution has been adjusted to obtain an acceptance rate close to 0.5, as recommended in [26]). Note finally that the independent positions (positions that are not directly related by the MRF-TV structure) are sampled in parallel using a check-board scheme, which accelerates the sampling procedure.

2) Sampling the reflectivity coefficients: Using (5) and (7), it can be easily shown that $\boldsymbol{r}$, and $\boldsymbol{w}$ are distributed according to the following gamma and inverse gamma distributions

$$r_{i,j}|\boldsymbol{\epsilon},\zeta \;\sim\; \mathcal{G}\left(4\zeta+c_2 k_{i,j} r_{i,j}^{\text{ML0}},\frac{1}{\beta_{i,j}}\right), \tag{15}$$

$$w_{i,j}|\boldsymbol{r},\zeta \;\sim\; \mathcal{IG}\left(4\zeta,4\zeta\rho_{1,i,j}(\boldsymbol{r})\right), \tag{16}$$

where $\beta_{i,j}=4\zeta\rho_{2,i,j}(\boldsymbol{w})+c_2 k_{i,j}\exp^{(-\alpha t_{i,j})}$, $k_{i,j}=0$ if the pixel is missing and $k_{i,j}=1$ otherwise (non-empty observed pixel). As a consequence, sampling according to (15) and (16) is straightforward.

3) Updating the MRF parameters: The MRF parameters maximizing the marginal likelihood $f\left(\boldsymbol{Y}|\eta,\zeta\right)$ are updated using the approach proposed in [14]. As reported in [3], [14], this approach provides a good approximation of the MRF parameters while requiring a reduced computational cost when compared to alternative approaches [25]. At each iteration of the MCMC algorithm, η and ζ are updated as follows

$$\eta^{(n+1)} = \mathcal{P}_{[0,\eta_{\max}]}\left\{\eta^{(n)} + \varsigma_n\left[TV\left(\boldsymbol{t}^{(n)}\right) - TV\left(\boldsymbol{t}'\right)\right]\right\} \tag{17}$$

$$\zeta^{(n+1)} = \mathcal{P}_{[0,\zeta_{\max}]}\left\{\zeta^{(n)} + \varsigma_n\left[\phi\left(\boldsymbol{r}^{(n)},\boldsymbol{w}^{(n)}\right) - \phi\left(\boldsymbol{r}',\boldsymbol{w}'\right)\right]\right\} \tag{18}$$

where $\varsigma_n = n^{-3/4}$, $\mathcal{P}_{[a,b]}(x)$ denotes the projection operator of x in the interval $[a,b]$, and $\phi\left(\boldsymbol{r},\boldsymbol{w}\right) = -4\sum_{(i,j)\in\nu_{\boldsymbol{w}}}\log\left(w_{i,j}\right) + 4\sum_{(i',j')\in\nu_{\boldsymbol{r}}}\log\left(r_{i',j'}\right) - \sum_{((i,j),(i',j'))\in\mathcal{E}}\left(\frac{r_{i',j'}}{w_{i,j}}\right)$. These expressions originate from a projected gradient descent step in which the intractable gradients $\frac{\partial}{\partial\eta}\log f\left(\boldsymbol{Y}|\eta^{(n)},\zeta^{(n)}\right)$ and $\frac{\partial}{\partial\zeta}\log f\left(\boldsymbol{Y}|\eta^{(n)},\zeta^{(n)}\right)$ have been approximated by the biased estimators $\left[TV\left(\boldsymbol{t}^{(n)}\right) - TV\left(\boldsymbol{t}'\right)\right]$ and $\left[\phi\left(\boldsymbol{r}^{(n)},\boldsymbol{w}^{(n)}\right) - \phi\left(\boldsymbol{r}',\boldsymbol{w}'\right)\right]$. These estimators use the current samples $\boldsymbol{t}^{(n)}$, $\boldsymbol{r}^{(n)}$, $\boldsymbol{w}^{(n)}$ and other auxiliary samples $\boldsymbol{t}',\boldsymbol{r}',\boldsymbol{w}'$ generated with kernels $\mathcal{K}_1$ and $\mathcal{K}_2$ whose target distributions are (6) and (7), respectively (see Algo. 1). Note also that the values obtained are projected using the operator $\mathcal{P}$ to guarantee the positivity constraints of η and ζ and the stability of the stochastic optimization algorithm ($\eta_{\max} = \zeta_{\max} = 20$ in the following). Note finally that the hyperparameters are only updated in the burn-in period ($n < N_{\text{bi}}$) and are fixed to their final values for the useful samples $N_{\text{bi}} \leq n \leq N_{\text{MC}}$. Finally we refer the reader to [3], [14] for more details regarding this procedure.

B. Optimization algorithm

This section describes an alternative to the MCMC algorithm which is based on a fast optimization algorithm. The latter maximizes the joint posterior (10) w.r.t. the parameters of interest to approximate the MAP estimator of $\boldsymbol{\Theta}$. The resulting optimization problem is tackled using a coordinate descent algorithm (CDA) [15], [16], [27] that sequentially updates the different parameters as illustrated in Algo. 2. Thus, the algorithm iteratively updates each parameter by maximizing its conditional distribution as described in the following subsections.

1) Updating the target positions: Maximizing the conditional distribution of the target positions (13) is equivalent to minimizing its negative logarithm $\mathcal{C}(\boldsymbol{t})$, given by (14). The latter is a proper, lower semi-continuous, coercive and strictly convex (since $r_{i,j}^{\text{ML0}} > 0, r_{i,j} > 0$)

Algorithm 1 MCMC algorithm

Input N_{bi}, N_{MC} and the impulse response parameters c_1, σ^2
Initialization
Initialize parameters $\boldsymbol{t}^{(0)}, \boldsymbol{r}^{(0)}, \boldsymbol{w}^{(0)}, \eta^{(0)}$, and $\zeta^{(0)}$
Update parameters/hyperparameters
for $n = 1 : N_{\text{MC}}$ **do**
  Sample $\boldsymbol{t}^{(n)}$ according to (13) using MH
  Sample $\boldsymbol{r}^{(n)}$ according to (15)
  Sample $\boldsymbol{w}^{(n)}$ according to (16)
  if $n < N_{\text{bi}}$ **then**
    Sample $\boldsymbol{t}' \sim \mathcal{K}_1(\boldsymbol{t}|\boldsymbol{t}^{(n)}, \eta^{(n-1)})$
    Sample $(\boldsymbol{r}', \boldsymbol{w}') \sim \mathcal{K}_2(\boldsymbol{r}, \boldsymbol{w}|\boldsymbol{r}^{(n)}, \boldsymbol{w}^{(n)}, \zeta^{(n-1)})$
    Update η using (17)
    Update ζ using (18)
  end if
end for
Output $\left\{\boldsymbol{t}^{(n)}, \boldsymbol{r}^{(n)}\right\}_{n=1}^{N_{\text{MC}}}$

function w.r.t. $\boldsymbol{t}$, so that there exists a unique minimizer of $\mathcal{C}(\boldsymbol{t})$ (see the Appendix).This problem can be solved using many convex programing algorithms [17], [18], [28], [29]. In this paper, we consider the ADMM variant proposed in [11] that has shown good performance in many fields [9], [30] while requiring a reduced computational cost. This algorithm is theoretically ensured to reach the unique minimum of $\mathcal{C}(\boldsymbol{t})$. More details regarding this algorithm and its convergence properties are provided in the Appendix.

2) Updating the reflectivity coefficients: Similarly to the target positions, maximizing the conditional distribution of $\boldsymbol{r}$ (resp. $\boldsymbol{w}$) provided in (15) (resp. (16)) is equivalent to minimizing $\mathcal{C}_1$ (resp. $\mathcal{C}_2$) given by

$$\mathcal{C}_1(\boldsymbol{r}) = \sum_{i,j} (1 - 4\zeta - c_2 k_{i,j} r_{i,j}^{\text{ML0}}) \log(r_{i,j}) + \frac{r_{i,j}}{\beta_{i,j}} \tag{19}$$

$$\mathcal{C}_2(\boldsymbol{w}) = \sum_{i,j} (4\zeta + 1) \log(w_{i,j}) + \frac{4\zeta \rho_{1,i,j}(\boldsymbol{r})}{w_{i,j}}. \tag{20}$$

The minimum of these functions is uniquely attained and given by

$$\overline{r_{i,j}} = \frac{4\zeta + c_2 k_{i,j} r_{i,j}^{\text{ML0}} - 1}{\beta_{i,j}}, \forall i,j \tag{21}$$

$$\overline{w_{i,j}} = \frac{4\zeta \rho_{1,i,j}(\boldsymbol{r})}{4\zeta + 1}, \forall i,j \tag{22}$$

subject to $4\zeta + c_2 r_{i,j}^{\text{ML0}} > 1$ which is always satisfied for $\zeta > 0.25$. These solutions are used to update the parameters $\boldsymbol{r}$ and $\boldsymbol{w}$ as shown in Algo. 2.

3) Convergence and stopping criteria: The proposition 2.7.1 in [15] asserts that the limit points of the sequence generated by the coordinate descent algorithm ($\boldsymbol{\Theta}^n$ for the nth iteration) are stationary points of $\mathcal{F} = -\log[f(\boldsymbol{\Theta}|\boldsymbol{Y}, \eta, \zeta)]$ provided that the minimum of that function w.r.t. $\boldsymbol{\Theta}$ along each coordinate is unique and that the function $\mathcal{F}$ is monotonically non-increasing along each coordinate in the interval from $\boldsymbol{\theta}_i^n$ to $\boldsymbol{\theta}_i^{n+1}$. These conditions are satisfied for the parameters considered. Indeed, the estimation of the target positions is a convex minimization problem whose solution is uniquely attained by the ADMM algorithm. Along the reflectivity coordinate, the function $\mathcal{C}_1$ is convex and has a unique minimum (for $\zeta > 0.25$). Along the auxiliary variable coordinate, $\mathcal{C}_2$ has a unique minimum and is monotonically non-increasing on each side of the minimum. These satisfy the conditions of the proposition 2.7.1 in [15]. Moreover, note that the cost function $\mathcal{F}$ is not convex, thus, the solution obtained might depend on the initial values that should be chosen carefully. Therefore, the reflectivity and target positions are initialized using the result of the classical approach (known as X-corr algorithm [3]). For each pixel, this approach estimates the reflectivity by $r_{i,j}^{\text{ML0}}$ and the depth by finding the maximum of the cross-correlation of the histogram $\boldsymbol{y}_{i,j}$ with the impulse response g_0 (see [3] for more details regarding the X-corr algorithm). With these initializations, the proposed algorithm reached minima of "good quality" in the considered simulations (see Sections V and VI).

Two stopping criteria have been considered for Algo. 2. The first criterion compares the new value of the cost function to the previous one and stops the algorithm if the relative error between these two values is smaller than a given threshold, i.e.,

$$|\mathcal{F}\left(\boldsymbol{\Theta}^{t+1}\right) - \mathcal{F}\left(\boldsymbol{\Theta}^{t}\right)| \leq \delta \mathcal{F}\left(\boldsymbol{\Theta}^{t}\right), \tag{23}$$

where $|.|$ denotes the absolute value. The second criterion is based on a maximum number of iterations N_{max}. These values have been fixed empirically to $(\delta, N_{\text{max}}) = (10^{-2}, 500)$ in the rest of the paper.

Algorithm 2 Coordinate descent algorithm (CDA)

Input $N_{\text{max}}, c_1, \sigma^2, \eta, \zeta$

Initialization

Initialize parameters $\boldsymbol{t}^{(0)}, \boldsymbol{r}^{(0)}, \boldsymbol{w}^{(0)}$ and $n \leftarrow 1$

conv$\leftarrow 0$,

Parameter update

while conv$= 0$ **do**

  Update $\boldsymbol{t}^{(n)}$ using Algo. 3

  Update $\boldsymbol{r}^{(n)}$ according to (21)

  Update $\boldsymbol{w}^{(n)}$ according to (22)

  Set conv$\leftarrow 1$ if the convergence criteria are satisfied

  $n \leftarrow n + 1$

end while

V. Simulation on synthetic data

This section evaluates the performance of the proposed algorithms on synthetic data with a known ground truth. All simulations have been implemented using MATLAB R2015a on a computer with Intel(R) Core(TM) i7- 4790 CPU@3.60GHz and 32GB RAM. The section is divided into two parts whose objectives are: 1) introducing the criteria used for the evaluation of the estimation results, and 2) analysis of the algorithms performance for different background levels.

A. Evaluation criteria

The restoration quality was evaluated qualitatively by visual inspection and quantitatively using the signal-to-reconstruction error ratio, SRE $= 10\log_{10}\left(\frac{||\boldsymbol{x}||^2}{||\boldsymbol{x}-\widehat{\boldsymbol{x}}||^2}\right)$, where $\boldsymbol{x}$ is the reference depth or reflectivity image , $\widehat{\boldsymbol{x}}$ is the restored image and $||\boldsymbol{x}||^2$ denotes the ℓ_2 norm given by $\boldsymbol{x}^T\boldsymbol{x}$. The returned values of this criterion are in decibel, the higher the better. The reference images are known for synthetic images. For real data, the estimated images with the MCMC approach in clear water, and with the highest acquisition time are considered as reference maps. As a result of the assumption of the absence of background photons, the proposed algorithms may be biased in a highly scattering environment. This effect is evaluated by considering the normalized-bias criterion given by N-Bias $= \frac{|\mathbb{E}[\boldsymbol{x}-\widehat{\boldsymbol{x}}]|}{|\mathbb{E}[\boldsymbol{x}]|}$.

We also provide some measures that are used in the experimental sections. We define one attenuation length (AL) as the distance after which the transmitted light power is reduced to $1/e$ of its initial value. If a target is located at range d from the sensor, its stand-off distance expressed in AL can be computed as AL $= \alpha d$. This measure is commonly used to highlight the attenuation affecting a given target [4], and will be considered when processing real data. Similarly to [31], we consider two other measures related to the background level. The first is the signal-to-background ratio given by SBR $= \frac{rc_1}{b}$. The second is the signal-to-noise ratio given by SNR $= \frac{rc_1}{\sqrt{rc_1+b}}$.

B. Effect of the background

In a highly scattering environment or with reduced acquisition times, the background level might increase w.r.t. the useful signal. This section evaluates this effect when considering synthetic (computer-simulated) data. A synthetic data cube has been generated according to model (1) with the following parameters $\alpha = 0, c_1 = 1000, \sigma^2 = 100, b_{i,j} = 1, \forall i, j$, $N_r = 100$ pixels, $N_c = 100$ pixels, and $T = 2000$ time bins where a time bin represents 2 picoseconds. The depth distance d corresponding to T bins can be computed as follows $d = \frac{Tc}{2n_e}$, where c is the speed of light and n_e is the refractive index of the propagation environment ($n_e = 1$ for the air and $n_e = 1.33$ for water). The synthetic data contains ten depths in the range $[12, 48]$ cm and ten reflectivity levels in the interval $r_{i,j} \in [0, 1]$, as shown in Fig. 4. The DR images are estimated using the proposed MCMC and CDA algorithms. The CDA algorithm requires the regularization parameters to be set manually. In this study, we provide the best performance (in terms of SRE) of this algorithm when testing the following values $\eta \in [0.01, 0.1, 0.5, 1, 2, 5]$ and $\zeta \in [0.3, 5, 10]$. The performance analysis is conduced w.r.t. the SBR criterion that evaluates the ratio between the useful signal levels $r_{i,j}c_1$ (whose variation depend on the reflectivity levels shown in Fig. 4) and the background levels $b_{i,j} = 1, \forall i, j$. Fig. 5 shows the obtained SRE for depth and reflectivity w.r.t. SBR. Overall, the proposed algorithms provide similar performance. For both depth and reflectivity, the figure shows a decreasing performance when the SBR ratio decreases. However, the depth SRE remains high even for SBR $= 1$. The reflectivity performance decreases log-linearly w.r.t. the SBR ratio and attains low SRE values for SBR $= 1$. This is mainly due to a reflectivity estimation bias in the presence of a high background level. Fig. 6 highlights this behavior and shows the estimation bias for depth and reflectivity. While the depth bias is always lower than 10%, the reflectivity shows high biases for low SBR $= 1$ which explains the low SRE

values. This bias can be corrected when processing real data using a look-up-table, however, this is beyond the scope of this paper. These results highlight the sensitivity of the estimated reflectivity to the background level while they confirm the good estimation of the depth image even for low SBR.

VI. Simulation using real data

This section evaluates the performance of the proposed restoration algorithms by conducting two experiments. In both cases, the targets were put underwater while varying the concentration of Maalox[2] to change the attenuation level (i.e., attenuation factor α) of the environment. The images were acquired in June 2016 in the laboratory at Heriot-Watt University, using a time-of-flight scanning sensor, based on TCSPC. The transceiver system and data acquisition hardware used for this work are broadly similar to that described in [4]. The overall system had a jitter of ≈ 60ps full width at half-maximum (FWHM) while we describe the other main parameters in Table I. The section is divided into three main parts. The first part highlights the reconstruction of the reflectivity obtained in the highly attenuating environment. The second part evaluates the restoration performance of the proposed algorithms while varying α. The third part studies the restoration limits of the proposed algorithms while varying both α and the acquisition time per pixel t_{acq}.

A. Restoration of the reflectivity level

It is clear from (2) that if two objects are located in a attenuating environment (defined by α) at a different distance from the sensor, they will be attenuated differently. This leads to the reflectivity distortion effect that is highlighted in this section. The experiment considers two reference targets (spectralon panels) with known reflectance (10 $\%$ and 99 $\%$), that are put inside a tank of water (dim. $40 \times 25 \times 25$cm). The 99% reflectance spectralon panel is located at a longer distance from the sensor than the one at 10%, as shown in Fig. 7. Five data cubes (with 150×150 pixels and 500 time bins) were acquired for different attenuation levels $\alpha \in [0.6, 5.2, 11.3, 14.8, 17.3]$ (obtained by varying the amount of Maalox in water). Fig. 8 shows the reflectivity images estimated by the classical and the proposed algorithms. For clear water $\alpha = 0.6$, the images show two levels of reflectivity related to the two spectralon

[2]Maalox is a commercially available antacid medicine that strongly affects scattering without inducing significant optical absorption.

panels, and separated by the edge of the spectralon which appears as blue vertical columns in the reflectivity maps. However, as α increases, the reflectivity levels of the classical algorithm decrease differently in the two regions, until we obtain a uniform reflectivity map (same level in the two regions) for $\alpha = 14.8$. Indeed, the return from the 99 $\%$ reflectance spectralon panel is attenuated more than the 10 $\%$ reflectance one, since it is located at a longer distance. This distortion effect is corrected by the proposed CDA and MCMC algorithms that recover the true reflectivity level under the different conditions of attenuation, as shown in Fig. 8 (middle) and (bottom). Fig. 9 shows the average of the rows of the reflectivity maps when varying α, for the three algorithms. When increasing α, the classical algorithm (red lines) presents decreasing levels that end-up to be the same for $\alpha = 14.8$ and slightly inversed for $\alpha = 17.3$. The CDA and MCMC algorithms provide almost the same reflectivity results under different levels of α. The observed small differences are mainly due to the presence of a high background noise for large α, which affects the restoration performance of the proposed algorithms.

B. Restoration of underwater depth and reflectivity images

This section evaluates the performance of the proposed restoration algorithms when considering six real data cubes (of size 120×120 pixels and 300 time bins) of a plastic pipe, put at a stand-off distance of 1.68m in water. Fig. 10 presents the experimental scheme and shows a picture of the plastic pipe target. The scans were performed with an acquisition time of 100ms per pixel and different attenuation levels as shown in Table II. The latter also shows the SBR and SNR levels estimated experimentally using a spectralon with known reflectivity. We provide these levels to link the analysis of this part to that on synthetic data.

Table III shows the SRE obtained with the algorithms. The algorithms proposed in this paper outperform the classical approach except for the reflectivity at the highest AL. In addition, the proposed algorithms show similar performance with slightly better results for CDA whose hyperparameters have been adjusted to provide the highest SRE. Note that the MCMC algorithm also provide good results while automatically adjusting the MRF hyperparameters. However, this is achieved at the cost of significantly longer processing time, as highlighted in Table IV. Figs. 11 and 12 show examples of the obtained depth and reflectivity images with the algorithms for different ALs. The depths are restored well

by the two algorithms while it can be seen that CDA over-smooths the pipe. The MCMC algorithm preserves more of the pipe contours while retaining some noise. These effects are mainly related to the estimated MRF hyperparameters that are different for the two algorithms. Considering the reflectivity images, the classical approach is largely affected by the environmental attenuation factor while the proposed algorithms obtain acceptable results for AL ≤ 7.5. For higher attenuation lengths, the restored reflectivity images are not satisfactory for two reasons: (i) the presence of a high background level and (ii) the measure of α is not too accurate because of the low signal level for these challenging scenarios, which affects the algorithms performance.

C. Performance w.r.t. the acquisition times and the attenuation factor

This section explores the performance of the proposed algorithms when dealing with a reduced number of photons due to a reduced acquisition time or an attenuating environment. This evaluation is important to state the possible level of attenuation that can be dealt with the proposed algorithms. In this experiment, we will consider the data used in the previous section with $t_{\text{acq}} = 100$ms (see Fig. 10). Note however that the data format of timed events allows the construction of photon timing histograms associated with shorter acquisition times, after measurement, as the system records the time of arrival of each detected photon. Here, we evaluate our algorithms for acquisition times ranging from 0.01ms to 100ms per pixel. Table V reports the percentage of non-empty pixels w.r.t. t_{acq} and AL. As expected, this percentage is higher for high t_{acq} or low AL. Figs. 13 and 14 show the SRE as a function of t_{acq} for different attenuation lengths. First note that the MRF parameters of the CDA algorithm have been adjusted to provide the best SRE results, which explain why CDA outperforms MCMC in some cases. As expected, the algorithms performance generally decreases while reducing the acquisition times or increasing the attenuation levels. As AL increases, the algorithms require more acquisition time (i.e., more informative pixels) in order to obtain an acceptable performance. The latter are generally obtained for a percentage higher than 30% of non-empty pixels and AL ≤ 7.5. For example, when AL $= 7.5$, the CDA algorithm requires that $t_{\text{acq}} > 10$ms to reach a good performance both for depth and reflectivity. Therefore, given an attenuating environment defined by α, these results allow the setting of the required acquisition times to obtain a given level of accuracy.

VII. Conclusions

This paper introduced a hierarchical Bayesian model and two estimation algorithms for the restoration of depth and reflectivity obtained in the limit of very low photon counts and significant attenuation. The algorithms were designed to provide the single-photon community with useful, relatively fast, and practical tools for the image restoration. Using some assumptions, a new formulation was introduced leading to a log-concave likelihood that is only expressed using preliminary estimates of the DR images. The restoration of these two images was achieved by considering two MRF based prior distributions ensuring spatial correlation between the pixels. The resulting joint posterior distribution was used to approximate the Bayesian estimators. First, a Markov chain Monte Carlo procedure based on a Metropolis-within-Gibbs algorithm was used to sample the posterior of interest and to approximate the MMSE estimators of the unknown parameters using the generated samples. Second, a coordinate descent approach using an alternating direction method of multipliers algorithm was used to approximate the maximum a posteriori estimators. Both algorithms showed comparable performance while providing different characteristics, i.e., the MCMC algorithm was fully automatic while the CDA algorithm required a reduced computational time. Results on both synthetic and real data showed the ability of the proposed algorithms to correct the reflectivity distortion effect, and to restore the depth and reflectivity images obtained in highly attenuating environments. Future work includes relaxing some of the assumptions of this paper, which might lead to better performance at the price of a higher computational cost. Generalizing the algorithms to account for target with multiple depth returns [19], [22] is also an interesting issue which is worthy of investigation.

Appendix
ADMM algorithm

Consider the optimization problem

$$\underset{\boldsymbol{t}}{\operatorname{argmin}}\, \mathcal{C}\left(\boldsymbol{t}\right) = \underset{\boldsymbol{t}}{\operatorname{argmin}} \sum_{j=1}^{J} g_j\left(\boldsymbol{H}^{(j)}\boldsymbol{t}\right) \tag{24}$$

where $\boldsymbol{t} \in \mathbb{R}^{N\times 1}$, $g_j : \mathbb{R}^{p_j} \to \mathbb{R}$ are closed, proper, convex functions, and $\boldsymbol{H}^{(j)} \in \mathbb{R}^{p_j\times N}$ are arbitrary matrices. After denoting $\boldsymbol{u}^{(j)} = \boldsymbol{H}^{(j)}\boldsymbol{z} \in \mathbb{R}^{p_j}$ and introducing the auxiliary variable $\boldsymbol{d}^{(j)} \in \mathbb{R}^{p_j}$, the authors in [11], [17] introduced the ADMM variant summarized in Algo. 3 to solve (24). This algorithm converges when the matrix $\boldsymbol{M} = \left[\sum_{j=1}^{J}\left(\boldsymbol{H}^{(j)}\right)^{\top}\boldsymbol{H}^{(j)}\right]$ has

full rank, and the optimization problems in line 10 are solved exactly or if their sequences of errors are absolutely summable [17]. In our case, we have

$$g_1\left(u_{i,j}^{(1)}\right) = \frac{\left(u_{i,j}^{(1)} - t_{i,j}^{\text{ML0}} + \alpha\sigma^2\right)^2}{\frac{2\sigma^2}{c_2 r_{i,j}^{\text{ML0}}}} + c_2 r_{i,j} \exp^{\left(-\alpha u_{i,j}^{(1)}\right)},$$

$$g_2\left(\boldsymbol{u}^{(2)}\right) = \eta||\boldsymbol{u}^{(2)}||_1, \text{ and } \quad g_3\left(\boldsymbol{u}^{(3)}\right) = i_{\mathbb{R}_+}\left(\boldsymbol{u}^{(3)}\right), \tag{25}$$

where $\boldsymbol{H}^{(1)} = \boldsymbol{K}$ is a $Q \times N$ binary matrix that contains a single non-zero value (equals to 1) on each line to model the loss of some image pixels and Q is the number of non-empty pixels, $\boldsymbol{H}^{(2)}$ denotes the TV linear operator as described in [11], and $\boldsymbol{H}^{(3)} = \mathbb{I}_N$. These matrices lead to $\boldsymbol{M} = \mathbb{I}_N + \boldsymbol{K}^\top\boldsymbol{K} + \boldsymbol{H}^{(2)\top}\boldsymbol{H}^{(2)}$ which is a full rank matrix ($\boldsymbol{K}^\top\boldsymbol{K}$ is a diagonal matrix whose values equal 0 in the position of missing pixels and 1 otherwise). The updates of $\boldsymbol{u}^{(2)}, \boldsymbol{u}^{(3)}$ in line 10 of Algo. 3 are straightforward and lead to exact solutions. For $\boldsymbol{u}^{(1)}$, the optimization problem has been solved using few iterations of the Newton method [15]. Regarding the solution of (24), note that $g_1 + g_2 + \text{TV}$ is proper, coercive, lower semi-continuous, and strictly convex for $r_{i,j}^{\text{ML0}} > 0$, and $r_{i,j} > 0$ (which is satisfied). Since $\boldsymbol{K}$ is injective, we obtain that $\mathcal{C}(\boldsymbol{t}) = g_1(\boldsymbol{K}\boldsymbol{t}) + g_2(\boldsymbol{t}) + \eta\text{TV}(\boldsymbol{t})$ is proper, coercive, lower semi-continuous, and strictly convex, thus, there is a unique minimizer for $\mathcal{C}(\boldsymbol{t})$ (see for example [15], [17], [32]). The authors invite the reader to consult [11], [17], [18] for more details regarding the ADMM algorithm and its convergence characteristics.

REFERENCES

[1] M.-C. Amann, T. M. Bosch, M. Lescure, R. A. Myllylae, and M. Rioux, "Laser ranging: a critical review of unusual techniques for distance measurement," *Opt. Eng.*, vol. 40, Jan. 2001.

[2] G. Buller and A. Wallace, "Ranging and three-dimensional imaging using time-correlated single-photon counting and point-by-point acquisition," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 13, no. 4, pp. 1006–1015, July 2007.

[3] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, "Lidar waveform based analysis of depth images constructed using sparse single photon data," *IEEE Trans. Image Process.*, vol. 25, no. 5, pp. 1935–1946, Mar. 2015.

[4] A. Maccarone, A. McCarthy, X. Ren, R. E. Warburton, A. M. Wallace, J. Moffat, Y. Petillot, and G. S. Buller, "Underwater depth imaging using time-correlated single-photon counting," *Opt. Express*, vol. 23, no. 26, pp. 33 911–33 926, Dec 2015.

[5] A. McCarthy, X. Ren, A. D. Frera, N. R. Gemmell, N. J. Krichel, C. Scarcella, A. Ruggeri, A. Tosi, and G. S. Buller, "Kilometer-range depth imaging at 1550 nm wavelength using an InGaAs/InP single-photon avalanche diode detector," *Opt. Express*, vol. 21, no. 19, pp. 22 098–22 113, Sep 2013.

[6] A. Kirmani, D. Venkatraman, D. Shin, A. Colaço, F. N. C. Wong, J. H. Shapiro, and V. K. Goyal, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.

Algorithm 3 ADMM for depth estimation

Initialization

Initialize $\boldsymbol{u}_0^{(j)}, \boldsymbol{d}_0^{(j)}, \forall j,\ \mu$. Set $k \leftarrow 0$, conv$\leftarrow 0$

while conv$= 0$ **do**

  for j=1:J **do**

    $\xi_k^{(j)} \leftarrow \boldsymbol{u}_k^{(j)} + \boldsymbol{d}_k^{(j)}$,

  end for

  $\boldsymbol{t}_{k+1} \leftarrow \boldsymbol{M}^{-1} \sum_{j=1}^{J} \left(\boldsymbol{H}^{(j)}\right)^{\top} \xi_k^{(j)}$,

  for j=1:J **do**

    $\boldsymbol{v}_k^{(j)} \leftarrow \boldsymbol{H}^{(j)} \boldsymbol{t}_{k+1} - \boldsymbol{d}_k^{(j)}$,

    $\boldsymbol{u}_{k+1}^{(j)} \leftarrow \underset{\boldsymbol{m}}{\operatorname{argmin}} \frac{\mu}{2} ||\boldsymbol{m} - \boldsymbol{v}_k^{(j)}||^2 + g_j(\boldsymbol{m})$,

  end for

  for j=1:J **do**

    $\boldsymbol{d}_{k+1}^{(j)} \leftarrow \boldsymbol{d}_k^{(j)} - \left(\boldsymbol{H}^{(j)} \boldsymbol{t}_{k+1} - \boldsymbol{u}_{k+1}^{(j)}\right)$,

  end for

  $k = k + 1$

end while

[7] A. M. Wallace, J. Ye, N. Krichel, A. McCarthy, R. Collins, and G. S. Buller, "Full waveform analysis for long-range 3d imaging laser radar," *EURASIP Journal on Advances in Signal Processing*, vol. 2010, no. 1, p. 896708, Dec. 2010.

[8] D. Shin, A. Kirmani, V. Goyal, and J. Shapiro, "Computational 3D and reflectivity imaging with high photon efficiency," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct 2014, pp. 46–50.

[9] A. Halimi, Y. Altmann, A. McCarthy, X. Ren, R. Tobin, G. S. Buller, and S. McLaughlin, "Restoration of intensity and depth images constructed using sparse single-photon data," in *Proc. EUSIPCO*, 2016, to appear.

[10] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992.

[11] M.-D. Iordache, J. Bioucas-Dias, and A. Plaza, "Total variation spatial regularization for sparse hyperspectral unmixing," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 11, pp. 4484–4502, Nov. 2012.

[12] O. Dikmen and A. Cemgil, "Gamma markov random fields for audio source modeling," *IEEE Trans. Audio, Speech, Language Process.*, vol. 18, no. 3, pp. 589–601, March 2010.

[13] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*. New York: Springer-Verlag, 1999.

[14] M. Pereyra, N. Whiteley, C. Andrieu, and J. Y. Tourneret, "Maximum marginal likelihood estimation of the granularity coefficient of a potts-markov random field within an mcmc algorithm," in *Proc. IEEE-SP Workshop Stat. and Signal Processing*, June 2014, pp. 121–124.

[15] D. P. Bertsekas, *Nonlinear programming*. Belmont, Massachusetts: Athena Scientific, 1995.

[16] J. Sigurdsson, M. Ulfarsson, and J. Sveinsson, "Hyperspectral unmixing with l_q regularization," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 11, pp. 6793–6806, Nov. 2014.

[17] M. Figueiredo and J. Bioucas-Dias, "Restoration of poissonian images using alternating direction optimization," *IEEE Trans. Image Process.*, vol. 19, no. 12, pp. 3133–3145, Dec 2010.

[18] M. Afonso, J. Bioucas-Dias, and M. Figueiredo, "An augmented lagrangian approach to the constrained optimization formulation of imaging inverse problems," *IEEE Trans. Image Process.*, vol. 20, no. 3, pp. 681–695, March 2011.

[19] S. Hernandez-Marin, A. Wallace, and G. Gibson, "Bayesian analysis of Lidar signals with multiple returns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 12, pp. 2170–2180, Dec. 2007.

[20] M. Carlavan and L. Blanc-Feraud, "Sparse poisson noisy image deblurring," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1834–1846, April 2012.

[21] Y. Altmann, A. Wallace, and S. McLaughlin, "Spectral unmixing of multispectral lidar signals," *IEEE Trans. Signal Process.*, vol. 63, no. 20, pp. 5525–5534, Oct 2015.

[22] S. Hernandez-Marin, A. M. Wallace, and G. J. Gibson, "Multilayered 3d lidar image construction using spatial models in a bayesian framework," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 6, pp. 1028–1040, June 2008.

[23] Y. Altmann, M. Pereyra, and S. McLaughlin, "Bayesian nonlinear hyperspectral unmixing with spatial residual component analysis," *IEEE Trans. Comput. Imaging*, vol. 1, no. 3, pp. 174–185, Sept 2015.

[24] A. Halimi, P. Honeine, and J. M. Bioucas-Dias, "Hyperspectral unmixing in presence of endmember variability, nonlinearity or mismodelling effects," *IEEE Trans. Image Process.*, 2016, to appear.

[25] M. Pereyra, N. Dobigeon, H. Batatia, and J. Tourneret, "Estimating the granularity coefficient of a Potts-Markov random field within a Markov Chain Monte Carlo algorithm," *IEEE Trans. Image Process.*, vol. 22, no. 6, pp. 2385–2397, June 2013.

[26] C. P. Robert and D. Cellier, "Convergence control of MCMC algorithms," in *Discretization and MCMC Convergence Assessment*, C. P. Robert, Ed. New York: Springer Verlag, 1998, pp. 27–46.

[27] A. Halimi, C. Mailhes, J.-Y. Tourneret, and H. Snoussi, "Bayesian estimation of smooth altimetric parameters: Application to conventional and delay/Doppler altimetry," *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 4, pp. 2207–2219, Mar. 2016.

[28] P. L. Combettes and J.-C. Pesquet, "A proximal decomposition method for solving convex variational inverse problems," *Inverse Problems*, vol. 24, no. 6, pp. x+27–x+27, 2008.

[29] L. Boyd, Sand Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.

[30] J. Bioucas-Dias and M. Figueiredo, "Alternating direction algorithms for constrained sparse regression: Application to hyperspectral unmixing," in *Proc. IEEE GRSS Workshop on Hyperspectral Image and SIgnal Processing: Evolution in Remote Sensing (WHISPERS)*, June 2010, pp. 1–4.

[31] S. Pellegrini, G. S. Buller, J. M. Smith, A. M. Wallace, and S. Cova, "Laser-based distance measurement using picosecond resolution time-correlated single-photon counting," *Measurement Science and Technology*, vol. 11, no. 6, p. 712, 2000.

[32] P. L. Combettes and V. R. Wajs, "Signal recovery by proximal forward-backward splitting," *SIAM Jour. Multiscale Modeling & Simulation*, vol. 4, no. 4, pp. 1168–1200, 2005.

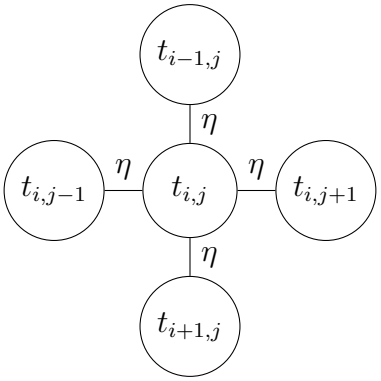

Fig. 1. The total variation neighborhood structure.

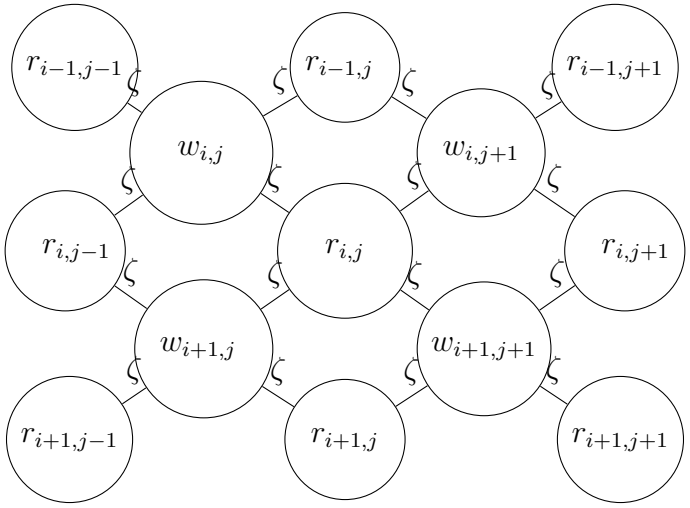

Fig. 2. Gamma-MRF neighborhood structure.

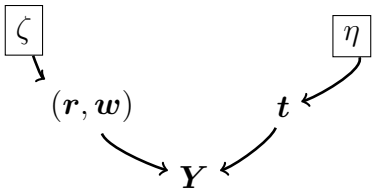

Fig. 3. DAG for the parameter and hyperparameter priors. For the optimization algorithm, the user fixed hyperparameters appear in boxes.

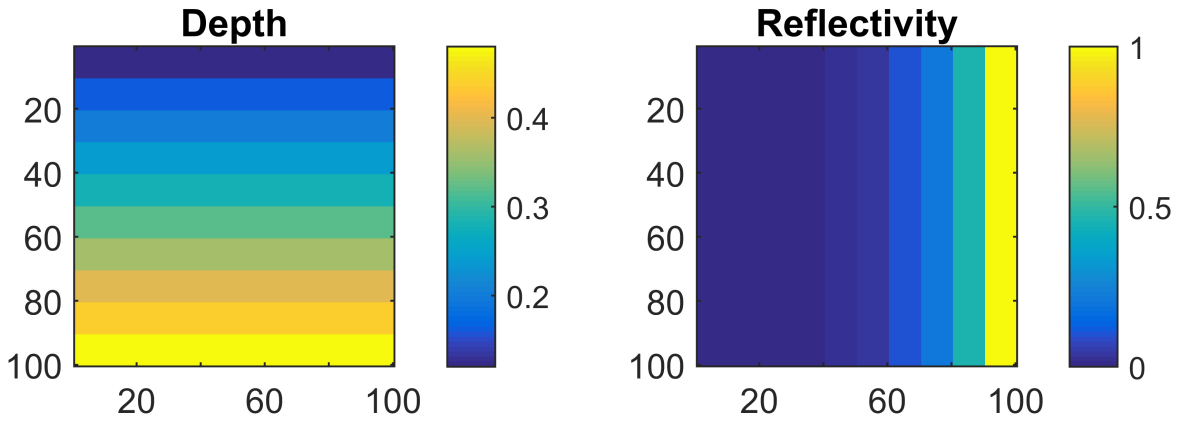

Fig. 4. Synthetic depth and reflectivity images.

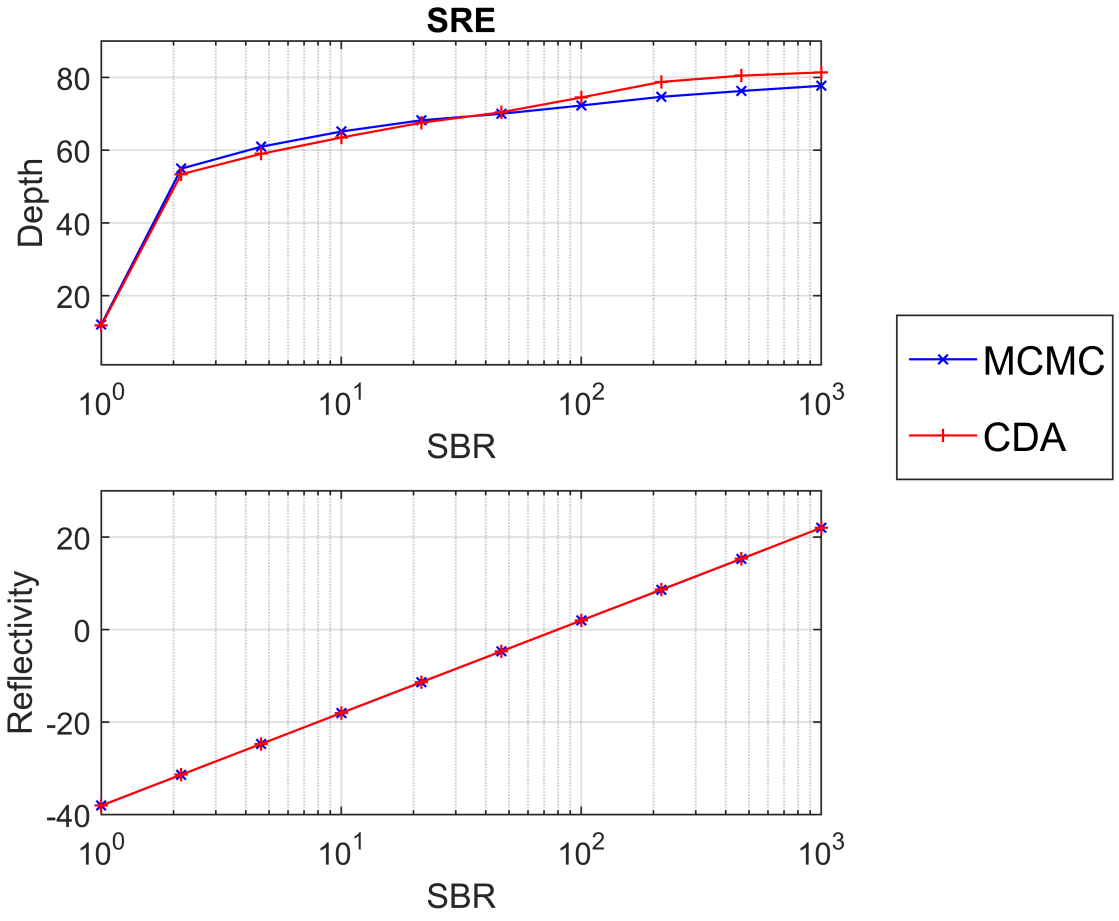

Fig. 5. SRE of depth and reflectivity with respect to the background levels for the MCMC (in blue) and CDA (in red) algorithms.

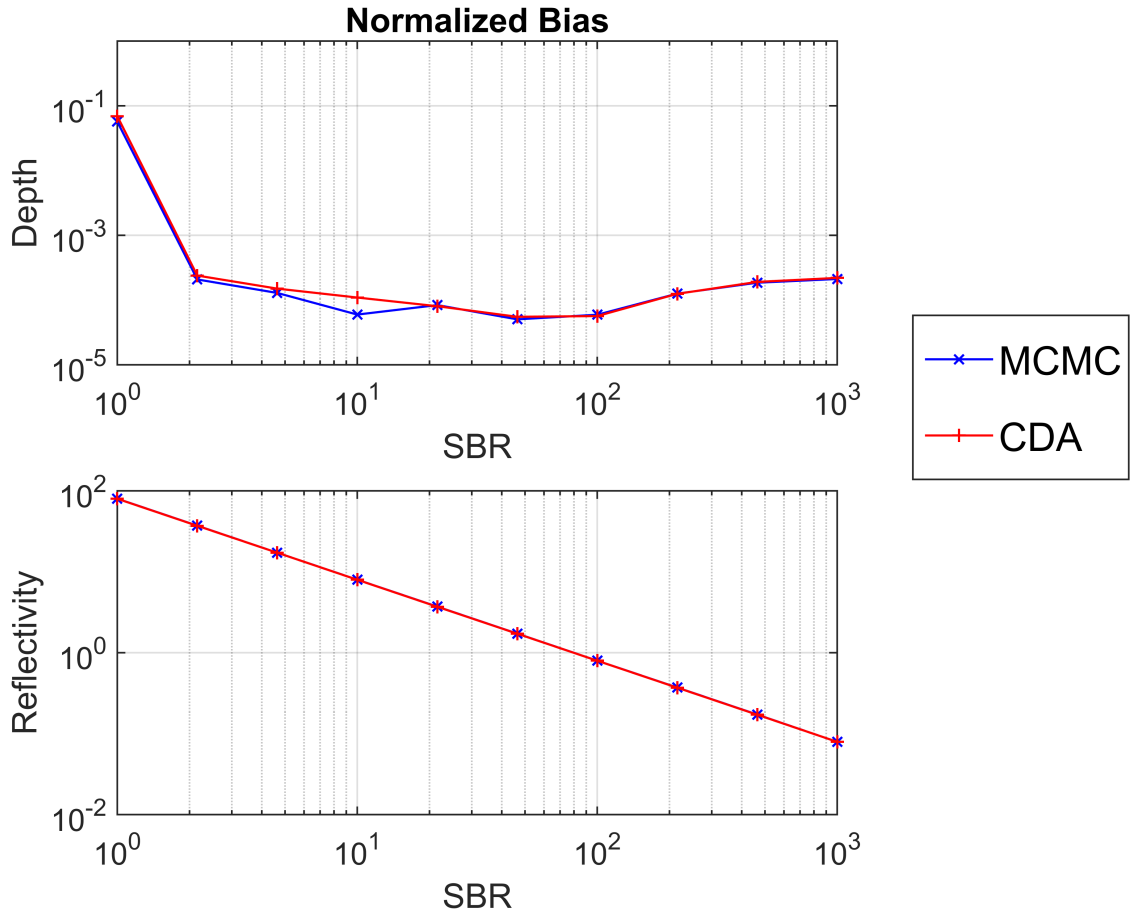

Fig. 6. Normalized bias of depth and reflectivity with respect to the background levels for the MCMC (in blue) and CDA (in red) algorithms.

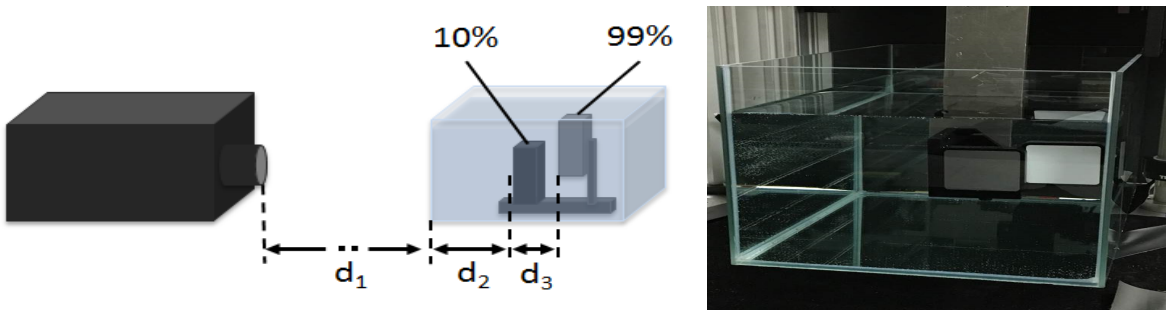

(a) (b)

Fig. 7. (a) Scheme of the first experiment with $d_1 = 1.57$m, $d_2 = 9.1$cm and $d_3 = 5.1$cm. (b) The two Spectralon targets.

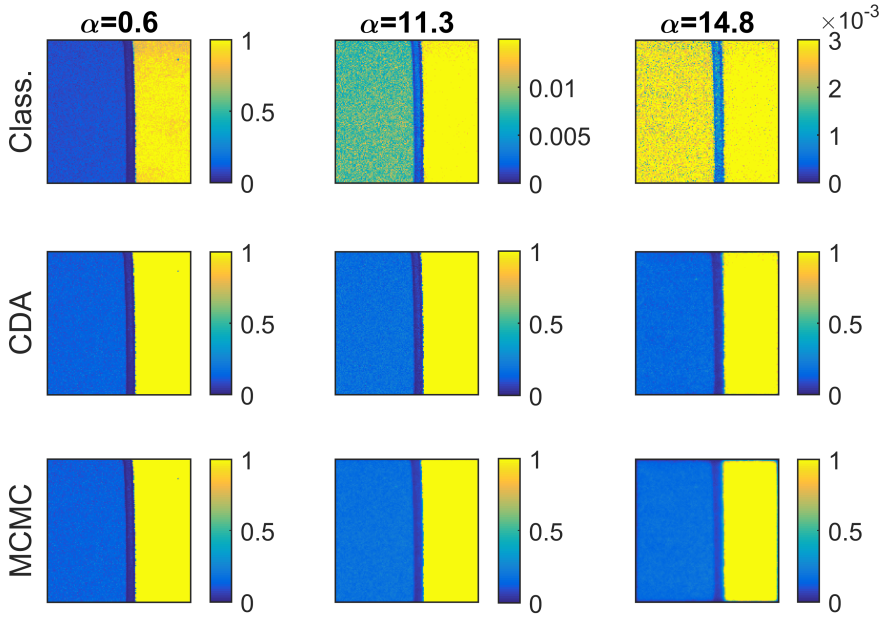

Fig. 8. Reflectivity images (150×150 pixels) obtained for $\alpha \in [0.6, 11.3, 14.8]$. (top) classical XCorr approach, (Middle) proposed CDA algorithm, (Bottom) proposed MCMC algorithm.

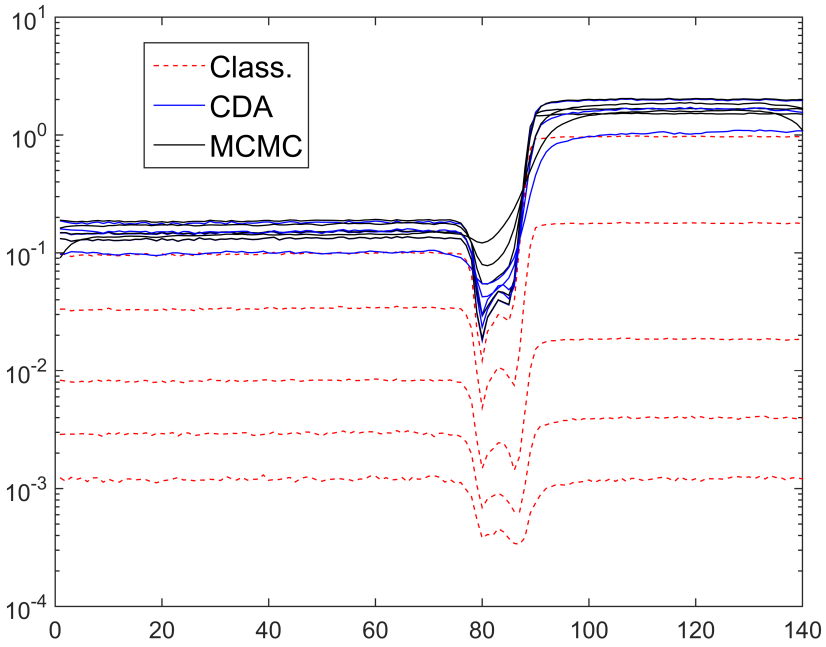

Fig. 9. Reflectivity lines (150 pixels) obtained for $\alpha \in [0.6, 5.2, 11.3, 14.8, 17.3]$ with the the classical XCorr approach (in dashed red lines), proposed CDA algorithm (in continuous blue lines) and the proposed MCMC algorithm (in continuous black lines).

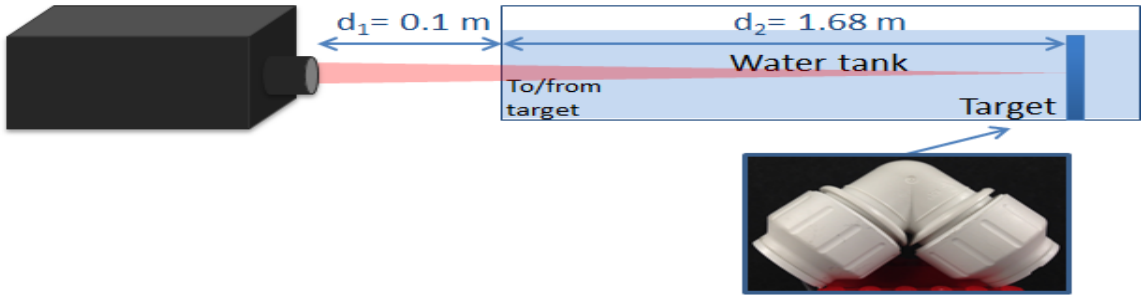

Fig. 10. Scheme of the second experiment showing a photograph of the plastic pipe target.

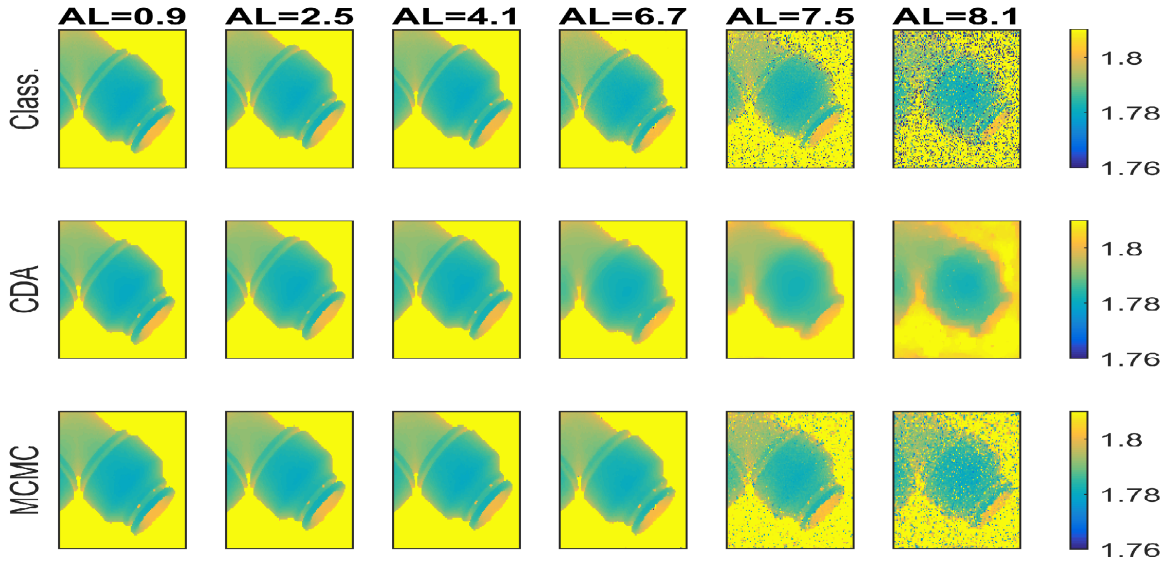

Fig. 11. Depth images (120×120 pixels) obtained for different attenuation factors with (top) the classical XCorr approach, (middle) the proposed CDA algorithm (bottom) and the proposed MCMC algorithm. The colormap is fixed for all images to [1.76,1.8] meters.

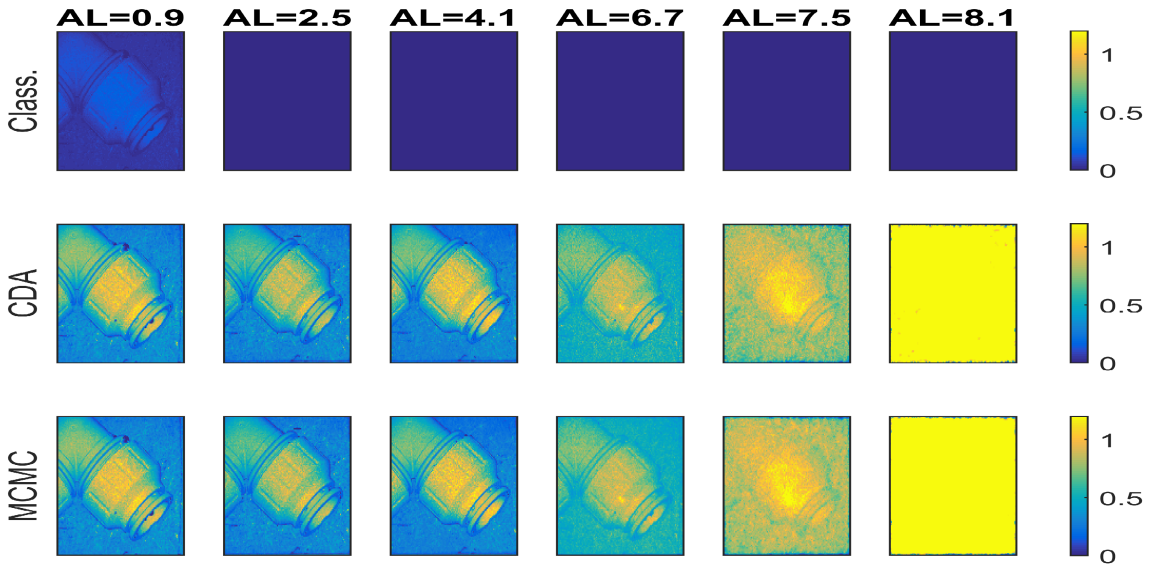

Fig. 12. Reflectivity images (120×120 pixels) obtained for different attenuation factors with (top) the classical XCorr approach, (middle) the proposed CDA algorithm (bottom) and the proposed MCMC algorithm. The colormap is fixed for all images to [0, 1.2].

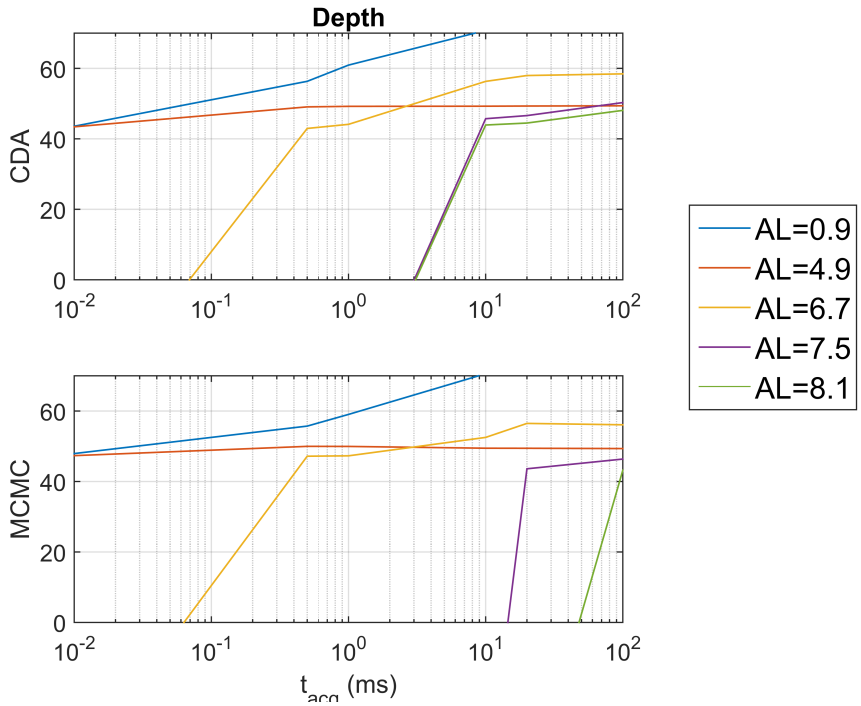

Fig. 13. Depth SRE obtained w.r.t. the acquisition time per-pixel (t_{acq}) for different attenuation levels. (top) CDA, (bottom) MCMC.

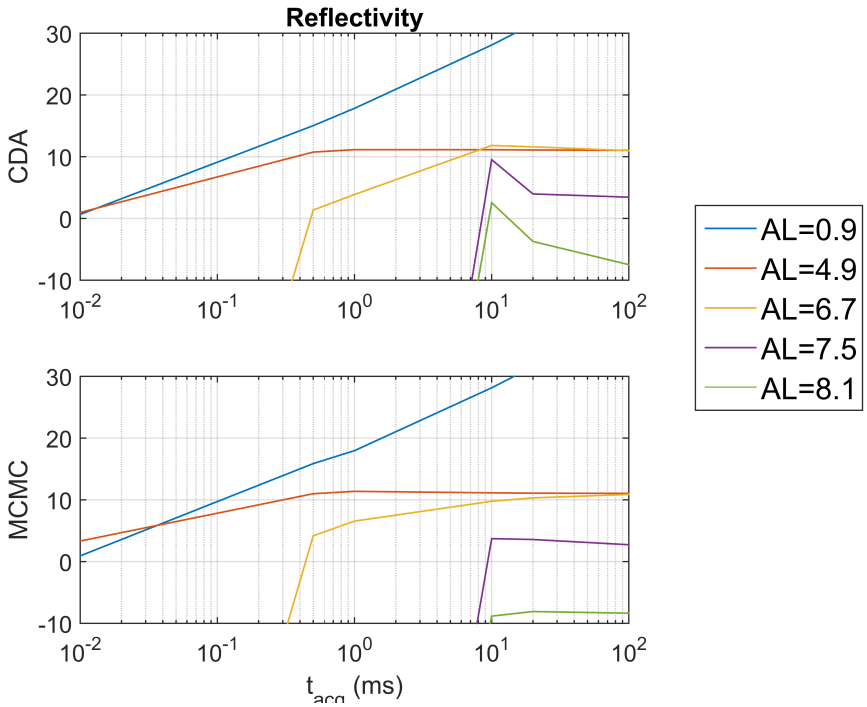

Fig. 14. Reflectivity SRE obtained w.r.t. the acquisition time per-pixel (t_{acq}) for different attenuation levels. (top) CDA, (bottom) MCMC.

TABLE I
MEASUREMENT PARAMETERS.

Laser system	Supercontinuum laser system
Illum. Wavelength	690nm
Laser Repetition Rate	19.5MHz
Histogram bin width	2ps

Target 1	2 reference targets with reflectivity 99% and 10% (see Fig. 7)
Scanned area	5×5cm
Number of pixels	150×150
Acquisition time	Per pixel: 10ms Total: ≈ 4 minutes
Histogram length	500bins (after gating)
Average optical power	$\approx$ 670nW

Target 2	Pipe ($\approx 8 \times 5 \times 3.5$cm) (see Fig. 10)
Scanned area	5×5cm
Number of pixels	120×120
Acquisition time	Per pixel: 100ms Total: ≈ 24 minutes
Histogram length	300bins (after gating)
Average optical power	see Table II

TABLE II
ATTENUATION LEVELS FOR THE UNDERWATER PIPE MEASUREMENT. THE CONCENTRATION OF MAALOX IS OBTAINED BY DIVIDING THE VOLUME OF MAALOX BY THE VOLUME OF THE WATER (67 LITERS).

Fraction of Maalox ($\times 10^{-4}$)	0	0.29	0.60	1	1.22	1.28
AL	0.9	2.5	4.1	6.7	7.5	8.1
SBR	2322	2576	2344	103	13	6
SNR	505	532	592	95	32	22
Average optical power (μW)	0.5	11	235	850	850	850

TABLE III
SRE (IN DB) OF THE RESTORED DEPTH AND REFLECTIVITY IMAGES W.R.T. THE ATTENUATION LENGTHS (AL).

		Attenuation lengths					
		0.9	2.5	4.1	6.7	7.5	8.1
Depth	Class.	71.7	49.9	49.3	54.8	36.4	34.2
	CDA	82.2	50.0	49.4	58.4	50.3	48.1
	MCMC	–	50.0	49.3	56.1	46.4	43.2
Reflectivity	Class.	1.5	0.0	0.0	0.0	0.0	0.0
	CDA	59.5	11.1	11.0	11.0	3.4	−7.5
	MCMC	–	11.1	11.0	10.9	2.7	−8.4

TABLE IV
PROCESSING TIME (IN SECONDS).

	Attenuation lengths					
	0.9	2.5	4.1	6.7	7.5	8.1
CDA	21	21	21	21	18	17
MCMC	529	513	514	524	496	494

TABLE V
PERCENTAGE OF USEFUL PIXELS W.R.T. t_{ACQ} AND AL.

		Attenuation lengths					
		0.9	2.5	4.1	6.7	7.5	8.1
t_{acq} (ms)	0.01	32.1	30.8	35.6	1.0	0.3	0.4
	0.1	91.9	91.2	92.1	9.2	3.3	2.5
	0.5	99.7	99.7	99.8	34.2	14.5	11.8
	1	99.9	100.0	100.0	51.5	25.0	20.7
	2	99.9	100.0	100.0	67.1	37.0	32.0
	10	100.0	100.0	100.0	88.6	53.5	43.0
	20	100.0	100.0	100.0	95.4	61.7	48.1
	100	100.0	100.0	100.0	100.0	84.0	72.6